A Parallel and Efficient Algorithm for Learning to Match

Jingbo Shang[1,4]*, Tianqi Chen[2]*, Hang Li[3], Zhengdong Lu[3], Yong Yu[4]

[1]University of Illinois at Urbana Champaign, IL, USA
shang7@illinois.edu

[2]University of Washington, Seattle, WA, USA
tqchen@cs.washington.edu

[3]Huawei Noah's Ark Lab, Hong Kong
{hangli.hl, lu.zhengdong}@huawei.com

[4]Shanghai Jiao Tong University, Shanghai, China
{shangjingbo, yyu}@apex.sjtu.edu.cn

arXiv:1410.6414v1 [cs.LG] 22 Oct 2014

ABSTRACT

Many tasks in data mining and related fields can be formalized as matching between objects in two heterogeneous domains, including collaborative filtering, link prediction, image tagging, and web search. Machine learning techniques, referred to as learning-to-match in this paper, have been successfully applied to the problems. Among them, a class of state-of-the-art methods, named feature-based matrix factorization, formalize the task as an extension to matrix factorization by incorporating auxiliary features into the model. Unfortunately, making those algorithms scale to real world problems is challenging, and simple parallelization strategies fail due to the complex cross talking patterns between subtasks. In this paper, we tackle this challenge with a novel parallel and efficient algorithm for feature-based matrix factorization. Our algorithm, based on coordinate descent, can easily handle hundreds of millions of instances and features on a single machine. The key recipe of this algorithm is an iterative relaxation of the objective to facilitate parallel updates of parameters, with guaranteed convergence on minimizing the original objective function. Experimental results demonstrate that the proposed method is effective on a wide range of matching problems, with efficiency significantly improved upon the baselines while accuracy retained unchanged.

1. INTRODUCTION

Many application tasks can be formalized as matching between objects in two heterogeneous domains, in which the association between some objects and information on those objects are given. We refer to the objects from one domain as queries and those from the other as targets, with the distinction usually clear from the context. For example, in collaborative filtering, given some items, one manages to find the users who have best match to the items, by using the preference of some users on some items as well as the features of users and items. Another example is image tagging, in which one wants to associate tags (keywords) with images based on some tagged images as well as the features of tags and images. Recent years have observed a great success of employing machine learning techniques, referred to as learning-to-match in this paper, to solve the matching problems.

Among existing approaches, a family of factorization models that make use of feature spaces to encode additional information, stand out as state-of-the-art in matching tasks. Examples include factorization machines [17, 18], feature-based latent factor models for link prediction [3, 14], and regression-based latent factor models [1]. We refer to this class of methods as feature-based matrix factorization (FMF) in this paper. The basic idea of FMF is to formalize the task as extension to plain matrix factorization for incorporating the features of objects into the model. In this way, one can make a full use of available information in the task to improve the accuracies. In fact, FMF is the best performer on many real world matching tasks. In collaborative filtering, FMF models using user feedback [9, 10], attribute [1, 23], and content [3, 24] have outperformed other models including plain matrix factorization. In web search, FMF models for calculating matching scores (relevance) between queries and documents have significantly enhanced relevance ranking [25, 26]. FMF models have also been successfully employed in link prediction [14], and have been adopted by the champion teams in KDD Cup 2012 [5, 18].

The learning of the FMF model can be conducted with a coordinate descent algorithm or a stochastic gradient descent algorithm. Since a matching problem is usually of a very large scale, with hundreds of millions of objects and features or more, it can easily become hard for FMF to manage. It is therefore necessary to develop a parallel and efficient algorithm for FMF. This is exactly the problem we attempt to address in this paper.

Making FMF scalable and efficient is much more difficult than it appears, due to the following two challenges. First, training requires simultaneous access to all the features, and thus the existing techniques for parallelization of matrix fac-

*These authors contributed equally to this work. The short version of this paper [22] was published in ICDM 2014.

torization [8, 27, 29] are not directly applicable. Second, the computation complexity of the coordinate descent algorithm is still too high, and it can easily fail to run on a single machine when the scale of problem becomes large, calling for techniques to significantly accelerate the computation. By making use of repeating patterns, the least-squares and probit losses can be scaled up for coordinate descent [19], but it does not provide guarantee for any general convex loss functions. Existed parallel coordinate descent algorithms, such as [6] and [15], due to the complex feature dependencies, cannot be directly applied here. The Hogwild! [16] algorithm for parallel stochastic gradient descent can be applied here, but it is a generic algorithm and thus is still inefficient for FMF.

In this paper, we try to tackle the two challenges by developing a parallel and efficient algorithm tailored for learning-to-match. The algorithm, referred to as parallel and efficient algorithm for learning-to-match (PL2M), parallelizes and accelerates the coordinate descent algorithm through (1) iteratively relaxing the objective to facilitate parallel updates of parameters, and (2) avoiding repeated calculations caused by features. The main contributions of this paper are as follows.

- We propose the parallel and efficient algorithm for feature-based matrix factorization, which iteratively relaxes the objective for parallel updates of parameters, and neatly avoids repeated calculations caused by features, for any general convex loss functions.
- We theoretically prove the convergence of the proposed algorithms on minimizing the original objective function, which is further verified by our extensive experiments. The parallel algorithm can automatically adjust the rate of parallel updates according to the conditions in learning.
- We empirically demonstrate the effectiveness and efficiency of the proposed algorithm on four benchmark datasets. The parallel algorithm achieves nearly linear speedup and the proposed acceleration helps the parallel algorithm run about 5 times faster than the Hogwild! [16] algorithm on average, using 8 threads.

Given the importance of the FMF models and difficulty of their parallelization, the work in this paper represents a significant contribution to the study of learning to match. To our best knowledge, this is the first effort on the scalability of the general FMF models.

The rest of the paper is organized as follows. Section 2 gives a formal description of the generalized matrix factorization and Section 3 explains the efficient coordinate descent algorithm. Section 4 describes parallelization of the coordinate descent algorithm. Related work is introduced in Section 5. Experimental results are provided in Section 6. Finally, the paper is concluded in Section 7.

2. LEARNING TO MATCH

In this section, we give a formal definition of learning-to-match and a formulation of a feature-based matrix factorization. We also present our motivation of parallelizing this learning task.

2.1 Problem Formulation

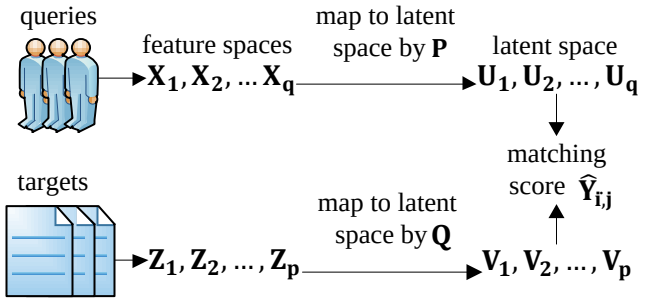

Figure 1: The general feature-based matrix factorization model for learning-to-match, which we have accelerated and parallelized in this paper.

Learning-to-match can be formally defined as follows. Let $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2 \cdots \mathbf{X}_q]$ be the instances in the query domain and $\mathbf{Z} = [\mathbf{Z}_1, \mathbf{Z}_2 \cdots \mathbf{Z}_p]$ be the instances in the target domain, where $\mathbf{X}_i \in \mathbb{R}^{n\times 1}$ and $\mathbf{Z}_j \in \mathbb{R}^{m\times 1}$ are query and target instances (feature vectors) respectively. For some query-target pairs, the corresponding matching scores $\{\mathbf{Y}_{ij}|(i,j) \in \mathcal{O}\}$ are given as training data, where $\mathcal{O}$ is the set of indices for all observed query-target pairs. Our problem is to learn to predict the matching score $\hat{\mathbf{Y}}_{ij} = f(\mathbf{X}_i, \mathbf{Z}_j)$ between any pair of query i and target j.

The setting is rather general and it subsumes many application problems. For example, in collaborative filtering, a user's preference over an item can be interpreted as the matching score between the user and the item. In social link prediction, the likelihood of link between nodes on the network can be as regarded as the matching score between the nodes. Web search, in general document retrieval, can also be formalized as a problem of first matching between a given query and documents and then ranking of documents based on the matching scores.

The goal of learning-to-match is to make accurate prediction by effectively using the information on the given relations between instances (e.g., similar users may prefer similar items), as well as the information on the features of instances (e.g., users may prefer items with similar properties).

2.2 Model

The query and target instances (feature vectors) are in two heterogeneous feature spaces, and a direct match between them is generally impossible. Instead, we map the feature vectors in the two domains into a latent space and perform matching on the images of the feature vectors in the latent space. We calculate the matching score of a query-target pair as

$$\hat{\mathbf{Y}}_{ij} = \mathbf{U}_i^T \mathbf{V}_j, \ \mathbf{U}_i = \mathbf{P}\mathbf{X}_i, \mathbf{V}_j = \mathbf{Q}\mathbf{Z}_j \tag{1}$$

where $\mathbf{P} \in \mathbb{R}^{d\times n}$ and $\mathbf{Q} \in \mathbb{R}^{d\times m}$ are transformation matrices that map feature vectors from the feature spaces into the latent space. $\mathbf{U}_i$ and $\mathbf{V}_j$ are latent factors of query instance i and target instance j. In this paper, we use $\mathbf{U}_i$ to denote the ith column of matrix $\mathbf{U}$ and $\mathbf{V}_j$ to denote the jth column of matrix $\mathbf{V}$. We refer to the model in Equation (1) as the model of feature-based matrix factorization. The model can be also interpreted as linear matching function of latent factors, in which the latent factor of each instance is also

Algorithm 1: Coordinate Descent Algorithm for Learning to Match

randomly initialize $\mathbf{P}, \mathbf{Q}$;
$\mathbf{U} \leftarrow \mathbf{PX}$, $\mathbf{V} \leftarrow \mathbf{QZ}$, $\hat{\mathbf{Y}} \leftarrow \mathbf{U}^T\mathbf{V}$ {calculate $\mathbf{U}$, $\mathbf{V}$, $\hat{\mathbf{Y}}$}
while *not converge* **do**
  for $k = 1$ ***to*** d **do**
    for $s = 1$ ***to*** n **do**
      $x \leftarrow \sum_{i,j \in \mathcal{O}} g_{ij} \mathbf{V}_{kj} \mathbf{X}_{is}$
      $y \leftarrow \beta \sum_{i,j \in \mathcal{O}} \mathbf{V}_{kj}^2 \mathbf{X}_{is}^2$
      $\Delta \mathbf{P}_{ks} \leftarrow T(x, y, \mathbf{P}_{ks}, \alpha, \lambda)$
      $\mathbf{P}_{ks} \leftarrow \mathbf{P}_{ks} + \Delta \mathbf{P}_{ks}$
    end
  end
  $\mathbf{U} = \mathbf{PX}$ {recalculate buffered $\mathbf{U}$}
  update $\mathbf{Q}$ in the same way as $\mathbf{P}$
end

linearly constructed from feature vectors. The query latent factor $\mathbf{U}_i$ can be expressed as $\mathbf{U}_i = \sum_{s=1}^{n} \mathbf{X}_{is} \mathbf{P}_s$. The target latent factor $\mathbf{V}_j$ can be expressed similarly.

The model, as shown in Figure 1, contains many existing models of feature-based matrix factorization as special cases [17, 18, 3, 14, 1]. When no "informative" features are available for objects of both domains, the feature matrices contain only the indices of the objects. Clearly, in such cases $\mathbf{X}$ and $\mathbf{Z}$ become identity matrices of sizes $q \times q$ and $p \times p$, the feature-based matrix factorization model naturally degenerates to the plain matrix factorization [11].

The objective of the learning task then becomes [1]:

$$\mathbf{P}, \mathbf{Q} = \text{argmin}_{\mathbf{P},\mathbf{Q}} \sum_{i,j \in \mathcal{O}} l(\hat{\mathbf{Y}}_{ij}, \mathbf{Y}_{ij}) + \Omega(\mathbf{P}) + \Omega(\mathbf{Q}) \quad (2)$$

Here $l(\hat{\mathbf{Y}}_{ij}, \mathbf{Y}_{ij})$ is a strongly convex loss function that measures the difference between the prediction $\hat{\mathbf{Y}}_{ij}$ and the ground truth $\mathbf{Y}_{ij}$. The loss function l can be square loss for regression

$$l(\hat{\mathbf{Y}}_{ij}, \mathbf{Y}_{ij}) = (\hat{\mathbf{Y}}_{ij} - \mathbf{Y}_{ij})^2 \quad (3)$$

or logistic loss for classification

$$l(\hat{\mathbf{Y}}_{ij}, \mathbf{Y}_{ij}) = \mathbf{Y}_{ij} \ln(1 + e^{-\hat{\mathbf{Y}}_{ij}}) + (1 - \mathbf{Y}_{ij}) \ln(1 + e^{\mathbf{Y}_{ij}}) \quad (4)$$

where $\mathbf{Y}_{ij} \in [0, 1]$. Ω is a regularization term based on elastic net [30] which includes l_2 (ridge) and l_1 (lasso) regularization as its special cases.

$$\Omega(\mathbf{P}) = \alpha ||\mathbf{P}||_1 + \frac{\lambda}{2} ||\mathbf{P}||_2^2, \ \Omega(\mathbf{Q}) = \alpha ||\mathbf{Q}||_1 + \frac{\lambda}{2} ||\mathbf{Q}||_2^2 \quad (5)$$

2.3 Performance

The use of features in learning-to-match is crucial for the accuracy of the task. Usually an FMF model, when properly optimized, can produce a higher accuracy in prediction than a model of plan matrix factorization (MF). This is because the FMF model can leverage more information for the prediction, particularly the feature information, while the MF model can only rely on relations between instances which are usually very sparse. For example, in the task of recommendation [3], only about 0.1% entries are observed. In Tencent Weibo dataset at KDD Cup 2012, about 70% of users in the test set have no following records in the training set [4]. As a result, MF cannot achieve satisfactory results in the tasks, while FMF models give the best results.

In fact, it has been observed the FMF models (with different types of features used) achieve state-of-the-art results on many different tasks, outperforming the models of MF with big margin. For example, in collaborative filtering, user feedback (SVD++) [9], user attribute [1], and product attribute [23] are incorporated into models to further improve the accuracies in prediction. In web search [25, 26], term vectors of queries and documents are used as features to significantly improve relevance ranking. FMF models also give the best results in link prediction in KDD Cup 2012 [14, 3, 4].

[1] $\mathbf{U}$ and $\mathbf{V}$ are not model parameters; they are only auxiliary variables determined by $\mathbf{P}$ and $\mathbf{Q}$.

Algorithm 2: Efficient Algorithm for Learning to Match

randomly initialize $\mathbf{P}, \mathbf{Q}$
$\mathbf{U} \leftarrow \mathbf{PX}$, $\mathbf{V} \leftarrow \mathbf{QZ}$, $\hat{\mathbf{Y}} \leftarrow \mathbf{U}^T\mathbf{V}$ {calculate $\mathbf{U}$, $\mathbf{V}$, $\hat{\mathbf{Y}}$}
while *not converge* **do**
  for $k = 1$ ***to*** d **do**
    for $i = 1$ ***to*** q **do**
      $G_{ki} \leftarrow \sum_{j \in \mathcal{O}_i} g_{ij} \mathbf{V}_{kj}$
      $H_{ki} \leftarrow \beta \sum_{j \in \mathcal{O}_i} \mathbf{V}_{kj}^2$
    end
    for $s = 1$ ***to*** n **do**
      $x \leftarrow \sum_i G_{ki} \mathbf{X}_{is}$
      $y \leftarrow \sum_i H_{ki} \mathbf{X}_{is}^2$
      $\Delta \mathbf{P}_{ks} \leftarrow T(x, y, \mathbf{P}_{ks}, \alpha, \lambda)$
      $\mathbf{P}_{ks} \leftarrow \mathbf{P}_{ks} + \Delta \mathbf{P}_{ks}$
      for $i = 1$ ***to*** q **do**
        $G_{ki} \leftarrow G_{ki} + \mathbf{X}_{is} \Delta \mathbf{P}_{ks} H_{ki}$
      end
    end
  end
  $\mathbf{U} = \mathbf{PX}$ {recalculate buffered $\mathbf{U}$}
  update $\mathbf{Q}$ in the same way as $\mathbf{P}$
end

2.4 Scalability

The success of the FMF models strongly indicates the necessity of scaling up the corresponding learning algorithms, given that the existing algorithms still cannot easily handle large datasets. By making use of repeating patterns, the least-squares and probit losses can be scaled up for coordinate descent [19], but it does not guarantee for any general convex loss functions. Other algorithms of parallel coordinate descent, such as [6] and [15] cannot be directly applied to FMF, because it is difficult for them to handle complex feature dependencies in FMF. The Hogwild! [16] algorithm for parallel stochastic gradient descent can be applied here, but it is a generic algorithm and thus is still inefficient for FMF. To our best knowledge, our work in this paper is the first effort on scalability of learning-to-match, i.e., feature-based matrix factorization.

3. EFFICIENT ALGORITHM FOR LEARNING TO MATCH

In this section, we propose an acceleration of the coordinate descent algorithm for solving the feature-based matrix factorization problem. We prove the convergence of the ac-

celerated algorithm, and we also give its time complexity at Section 4.3.

3.1 Coordinate Descent Algorithm

Let $g_{ij} = \frac{\partial}{\partial \hat{\mathbf{Y}}_{ij}} l(\hat{\mathbf{Y}}_{ij}, \mathbf{Y}_{ij})$ be the gradient of each instance over prediction and β be constant of l, such that,

$$l(y + dy) \leq l(y) + l'(y)dy + \frac{\beta}{2}dy^2 \tag{6}$$

That is, $\beta = \sup_y \frac{\partial^2}{\partial y^2} l(y, \mathbf{Y}_{ij})$. We can exploit the standard technique to learn the model using coordinate descent (CD), as shown in Algorithm 1[2]. Here, T is defined using the following thresholding function to handle optimization of l_1 norm

$$T(x, y, w, \alpha, \lambda) = \begin{cases} \max(-\frac{x+\lambda w+\alpha}{y+\lambda}, -w) & w - \frac{x+\lambda w}{y+\lambda} \geq 0 \\ \min(-\frac{x+\lambda w-\alpha}{y+\lambda}, -w) & w - \frac{x+\lambda w}{y+\lambda} < 0 \end{cases} \tag{7}$$

The regularization term only affects the result through function T, and thus makes most part of the algorithm independent of regularization. Note that we implicitly assume $\hat{\mathbf{Y}}_{ij}$ is buffered and kept up to date, when g_{ij} is needed in the algorithm.

The time complexity of one update in Algorithm 1 is $O(d|\mathcal{O}|(\frac{\|\mathbf{X}\|_0}{q} + \frac{\|\mathbf{Z}\|_0}{p}))$, where $\|\mathbf{X}\|_0$ and $\|\mathbf{Z}\|_0$ denote numbers of nonzero entries in feature matrices $\mathbf{X}$ and $\mathbf{Z}$, q and p denote numbers of query and target instances, and $(\frac{\|\mathbf{X}\|_0}{q} + \frac{\|\mathbf{Z}\|_0}{p})$ denotes average number of nonzero features for each pair $(i, j) \in \mathcal{O}$. We note that the time complexity is the same as the time complexity of stochastic gradient optimization for Equation (2). From the analysis, we can see that the time complexity of algorithm is increased by order of $O(d|\mathcal{O}|)$ when average number of nonzero features increases. This can greatly hamper the learning of matching model, when a large number of features are used.

3.2 Acceleration

We give an efficient algorithm for learning-to-match by avoiding the repeated calculations caused by features. There is only a little works focused on the acceleration of the FMF models. The most relevant one is that scaling up some specific coordinate descent by making use of repeating patterns of features [19]. However, it is specialized for the least-squares and probit losses. Although the idea is similar to the avoiding repeated calculations in our work, we extend the idea to any general convex loss functions.

One can see that there exist repeated calculations of summations for the same query or target when calculating x and y in Algorithm 1, which gives us a chance to speed up the algorithm. We introduce two auxiliary variables G_{ki} and H_{ki} calculated by

$$G_{ki} = \sum_{j \in \mathcal{O}_i} g_{ij}\mathbf{V}_{kj}, \; H_{ki} = \beta \sum_{j \in \mathcal{O}_i} \mathbf{V}_{kj}^2 \tag{8}$$

where $\mathcal{O}_i = \{j | (i, j) \in \mathcal{O}\}$ is the set of observed target instances associated with query instance i. The key idea of efficient CD is to make use of G_{ki} and H_{ki} to save duplicated summations in Algorithm 1. Since the gradient value g_{ij} is changed after each update, it is not trivial to let G_{ki} unchanged. Our algorithm keeps making G_{ki} updated to ensure the convergence of the algorithm. The efficient algorithm for learning-to-match is shown in Algorithm 2.

[2]In this paper, all matrix operations in algorithms have taken the advantages of sparsity, e.g. summations are all over nonzero entries implicitly. So do the time complex analysis and implementations of all algorithms.

3.3 Convergence of Algorithm

Next, we prove the convergence of Algorithm 2, which greatly reduces the time complexity of learning. Suppose that the kth row of $\mathbf{P}$ is changed by $\Delta\mathbf{P}_{k,:}$. After the change, the loss function can be bounded by

$$\begin{aligned} L(\Delta\mathbf{P}_{k,:}) &= \sum_{i,j \in \mathcal{O}} l(\hat{\mathbf{Y}}_{ij} + (\sum_s \Delta\mathbf{P}_{ks}\mathbf{X}_{is})\mathbf{V}_{kj}, \mathbf{Y}_{ij}) \\ &\leq \sum_{i,j \in \mathcal{O}} l(\hat{\mathbf{Y}}_{ij}, \mathbf{Y}_{ij}) + \sum_{i,j \in \mathcal{O}} g_{ij}\mathbf{V}_{kj}(\sum_s \Delta\mathbf{P}_{ks}\mathbf{X}_{is}) \\ &\quad + \frac{1}{2}\sum_{i,j \in \mathcal{O}} h_{ij}\mathbf{V}_{kj}^2(\sum_s \Delta\mathbf{P}_{ks}\mathbf{X}_{is})^2 \\ &= \sum_{i,j \in \mathcal{O}} l(\hat{\mathbf{Y}}_{ij}, \mathbf{Y}_{ij}) + \sum_i G_{ki}(\sum_s \Delta\mathbf{P}_{ks}\mathbf{X}_{is}) \\ &\quad + \frac{1}{2}\sum_i H_{ki}(\sum_s \Delta\mathbf{P}_{ks}\mathbf{X}_{is})^2 \\ &\triangleq \tilde{L}(\Delta\mathbf{P}_{k,:}) \end{aligned} \tag{9}$$

Intuitively, updating G_{ki} corresponds to minimizing the quadratic upper bound $\tilde{L}(\Delta\mathbf{P}_{k,:}) + \Omega(\mathbf{P}_{k,:} + \Delta\mathbf{P}_{k,:})$ of the original convex loss which is re-estimated each round. Formally, all the values of $\Delta\mathbf{P}_{k,:}$ are zero in the beginning. We need to sequentially update $\mathbf{P}_{ks}$ for different s's to minimize $\tilde{L}(\Delta\mathbf{P}_{k,:}) + \Omega(\mathbf{P}_{k,:} + \Delta\mathbf{P}_{k,:})$. Assuming that we have already updated $\mathbf{P}_{k,1} \cdots \mathbf{P}_{k,t-1}$ and need to decide $\Delta\mathbf{P}_{kt}$, we can calculate the upper bound $\tilde{L}$ as follows

$$\begin{aligned} \tilde{L} =& c + \sum_i \left(G_{ki} + \sum_{s=1}^{t-1} \mathbf{X}_{is}\Delta\mathbf{P}_{ks}H_{ki} \right) \mathbf{X}_{is}\Delta\mathbf{P}_{kt} \\ &+ \frac{1}{2}\left(\sum_i H_{ki}\mathbf{X}_{it}^2 \right) \Delta\mathbf{P}_{kt}^2 \end{aligned} \tag{10}$$

The first order term of this equation is exactly the update rule in Algorithm 2. Using $\Delta\mathbf{P}^*_{k,:}$ to denote the change on the kth row after carrying out the update, we arrive at the following inequality

$$\begin{aligned} L(\Delta\mathbf{P}^*_{k,:}) + \Omega(\mathbf{P}_{k,:} + \Delta\mathbf{P}^*_{k,:}) \leq& \tilde{L}(\Delta\mathbf{P}^*_{k,:}) + \Omega(\mathbf{P}_{k,:} + \Delta\mathbf{P}^*_{k,:}) \\ \leq& \tilde{L}(\mathbf{0}) + \Omega(\mathbf{P}_{k,:} + \mathbf{0}) \\ =& L(\mathbf{0}) + \Omega(\mathbf{P}_{k,:} + \mathbf{0}) \end{aligned} \tag{11}$$

Note that we start from $\Delta\mathbf{P}_{k,:} = \mathbf{0}$, and we have $\Delta\mathbf{P}_{k,:} = \Delta\mathbf{P}^*_{k,:}$ after the update. The inequality in Equation (11) shows that the original loss function decreases after each round of update, and hence this proves the convergence of Algorithm 2 for any differentiable convex loss function.

4. PARALLEL AND EFFICIENT ALGORITHM FOR LEARNING TO MATCH

We propose a parallel and efficient learning-to-match algorithm (PL2M) to further improve the scalability and efficiency by deriving an adaptive estimation of the conflicts

caused by parallel updates. Specifically, we consider parallelizing and accelerating Algorithm 2. The statistics calculation and preprocessing steps in Algorithm 2 can be naturally separated into several independent tasks and thus fully parallelized. However, there is strong dependency within $\mathbf{P}$ update steps, making the parallelization of it a difficult task. We will discuss how we solve the problem next.

4.1 Parallelization

Let S be a set of feature indices to be updated in parallel. Assume that the statistics of G_{ki} is up to date as in Algorithm 2 and we want to change $\mathbf{P}_{ks}$ for $s \in S$ in parallel. For simplicity of notation, we use $\Delta\mathbf{P}_{k,:}$ to represent the change in $\mathbf{P}_{k,:}$. The value of $\tilde{L}$ after this change will be

$$\begin{aligned}\tilde{L}(\Delta\mathbf{P}_{k,:}) = &\sum_{i,j\in\mathcal{O}} l(\hat{\mathbf{Y}}_{ij}, \mathbf{Y}_{ij}) \\ &+ \sum_{s\in S}(\sum_i G_{ki}\mathbf{X}_{is}\Delta\mathbf{P}_{ks} + \frac{1}{2}\sum_i H_{ki}\mathbf{X}_{is}^2\Delta\mathbf{P}_{ks}^2) \\ &+ \sum_{s\in S}\sum_{t\in S, t\neq s}\sum_i H_{ki}\mathbf{X}_{it}\mathbf{X}_{is}\Delta\mathbf{P}_{kt}\Delta\mathbf{P}_{ks}\end{aligned} \tag{12}$$

In a specific case in which $\mathbf{X}_{it}\mathbf{X}_{is} = 0$ for $s \neq t; s, t \in S$, the third line in Equation (12) becomes zero. This means that the features in the selected set do not appear in the same instance. In such case, the loss can be separated into $|S|$ independent parts and the original update rule can be applied in parallel. Not surprisingly, such condition does not hold in many real world scenarios. We need to remove these troublesome cross terms in the second line, by deriving an adaptive estimation of the conflicts caused by parallel updates, more specifically, by the inequality:

$$\Delta\mathbf{P}_{kt}\Delta\mathbf{P}_{ks}\mathbf{X}_{it}\mathbf{X}_{is} \leq \frac{1}{2}|\mathbf{X}_{it}\mathbf{X}_{is}|\left(\Delta\mathbf{P}_{kt}^2 + \Delta\mathbf{P}_{ks}^2\right) \tag{13}$$

With the inequality, we can bound $\tilde{L}$ as follows

$$\begin{aligned}\tilde{L}(\Delta\mathbf{P}_{k,:}) \leq &c + \sum_{s\in S}(\sum_i G_{ki}\mathbf{X}_{is}\Delta\mathbf{P}_{ks}) \\ &+ \sum_{s\in S}(\frac{1}{2}\sum_i H_{ki}|\mathbf{X}_{is}|(\sum_{t\in S}|\mathbf{X}_{it}|)\Delta\mathbf{P}_{ks}^2) \\ \triangleq &\tilde{L}^p(\Delta\mathbf{P}_{k,:}).\end{aligned} \tag{14}$$

Obviously this new upper bound $\tilde{L}^p(\Delta\mathbf{P}_{k,:})$ can be separated into $|S|$ independent parts and optimized in parallel. Moreover, the sum $\sum_{t\in S}|\mathbf{X}_{it}|$ is common for all features in set S and is only needed to be calculated once. With this result, we give a parallel and efficient algorithm for learning-to-match, shown in Algorithm 3.

4.2 Convergence of Algorithm

The relaxation of $\tilde{L}(\Delta\mathbf{P}_{k,:})$ into $\tilde{L}^p(\Delta\mathbf{P}_{k,:})$ is performed iteratively in the optimization, and it still attempts to optimize the original objective as in Equation (2), which is a case much analogous to Expectation-Maximization algorithm in finding a maximum-likelihood solution. Let $\Delta\mathbf{P}_{k,:}^*$ be the change in $\mathbf{P}_{k,:}$ after each parallel update. Since each parallel update optimizes $\tilde{L}^p(\Delta\mathbf{P}_{k,:})$, we have the following

Algorithm 3: Parallel Algorithm for Learning to Match

```
randomly initialize P, Q
U ← PX, V ← QZ, Ŷ ← U^T V {calculate U, V, Ŷ}
while not converge do
    schedule a partition P of {1, 2, ··· , n}
    for k = 1 to d do
        for i = 1 to q in parallel do
            G_ki ← Σ_{j∈O_i} g_ij V_kj
            H_ki ← β Σ_{j∈O_i} V_kj^2
        end
        for each index set S ∈ P do
            for i = 1 to q in parallel do
                C_ki ← Σ_{s∈S} |X_is|
            end
            for s ∈ S in parallel do
                x ← Σ_i G_ki X_is
                y ← Σ_i H_ki |X_is| C_ki
                ΔP_ks ← T(x, y, P_ks, α, λ)
                P_ks ← P_ks + ΔP_ks
            end
            for i = 1 to q in parallel do
                G_ki ← G_ki + Σ_{s∈S} X_is ΔP_ks H_ki
            end
        end
    end
    U = PX {recalculate buffered U}
    update Q in the same way as P
end
```

inequality

$$\begin{aligned}&\tilde{L}(\Delta\mathbf{P}_{k,:}^*) + \Omega(\mathbf{P}_{k,:} + \Delta\mathbf{P}_{k,:}^*) \\ \leq &\tilde{L}^p(\Delta\mathbf{P}_{k,:}^*) + \Omega(\mathbf{P}_{k,:} + \Delta\mathbf{P}_{k,:}^*) \\ \leq &\tilde{L}^p(\mathbf{0}) + \Omega(\mathbf{P}_{k,:} + \mathbf{0}) \\ = &\tilde{L}(\mathbf{0}) + \Omega(\mathbf{P}_{k,:} + \mathbf{0})\end{aligned} \tag{15}$$

It indicates that $\tilde{L}$ decreases after each parallel update. It then follows that the parallel procedure for optimizing the original loss function in Algorithm 3 always converges.

The update rule depends on the statistics $C_{ki} = \sum_{t\in S}|\mathbf{X}_{it}|$. With the following notation

$$\eta_{ks} = \frac{\sum_i H_{ki}\mathbf{X}_{is}^2 + \lambda}{\sum_i H_{ki}C_{ki}|\mathbf{X}_{is}| + \lambda} \tag{16}$$

It can be shown that the parallel update of $\Delta\mathbf{P}_{ks}$ is shrunken by η_{ks} compared to sequential update. Intuitively η_{ks} depends on the co-occurrence between features $s \in S$. When features in S rarely co-occur, η_{ks} will be close to one, which means that we can update "aggressively". When features in S co-occur frequently, η_{ks} will get small and we need to update more "conservatively". In an extreme case in which no feature co-occurs with each other, $\eta_{ks} = 1$ and we get perfect parallelization without any loss of update efficiency. In another extreme case in which we have $|S|$ duplicated features ($\mathbf{X}_{is} = \mathbf{X}_{it}$ for all $s, t \in S$), $\eta_{ks} = \frac{1}{|S|}$, which is extremely conservative given the size of S. The advantage of our algorithm is that it *automatically adjusts* its "level of conservativeness" by the condition in learning, and thus it always ensures the convergence of the algorithm regardless of the number of threads and the nature of dataset.

The changes in loss function can be analyzed accordingly. Let us consider the simple case in which $\alpha = 0$ and only l_2 regularization is involved. The change of loss after parallel update can be bounded by

$$\begin{aligned}
&\tilde{L}(\Delta\mathbf{P}^*_{k,:}) + \Omega(\mathbf{P}_{k,:} + \Delta\mathbf{P}^*_{k,:}) - \tilde{L}(\mathbf{0}) - \Omega(\mathbf{P}_{k,:} + \mathbf{0}) \\
\leq& -\frac{1}{4}\sum_{s\in S}\frac{\left(\sum_i G_{ki}\mathbf{X}_{is}\right)^2}{\sum_i H_{ki}|\mathbf{X}_{is}|C_{ki} + \lambda} \\
=& -\frac{1}{4}\sum_{s\in S}\eta_{ks}\frac{\left(\sum_i G_{ki}\mathbf{X}_{is}\right)^2}{\sum_i H_{ki}\mathbf{X}^2_{is} + \lambda}
\end{aligned} \tag{17}$$

As this inequality indicates, compared to the ideal case in which features do not co-occur, each parallel update's contribution to the loss change is scaled by η_{ks}. The above analysis also intuitively justifies that η_{ks} controls the efficiency of the update.

4.3 Time Complexity

The time complexity of the efficient algorithm (Algorithm 2) is only of $O(d(\|\mathbf{X}\|_0 + \|\mathbf{Z}\|_0 + |\mathcal{O}|))$. It is linear to numbers of nonzero entries of feature matrices and number of observed entries of $\mathbf{Y}$. Recall the time complexity of the coordinate descent algorithm (Algorithm 1), which is $O(d|\mathcal{O}|(\frac{\|\mathbf{X}\|_0}{q} + \frac{\|\mathbf{Z}\|_0}{p}))$.

The *speedup* on $\mathbf{P}$ updates in Algorithm 2 is as follows:

$$\frac{O(d|\mathcal{O}|\frac{\|\mathbf{X}\|_0}{q})}{O(d\|\mathbf{X}\|_0)} = O(\frac{|\mathcal{O}|}{q}) \tag{18}$$

This corresponds to average number of observed target instances per query instance. Similarly, on the $\mathbf{Q}$ updates, the *speedup* is about $O(\frac{|\mathcal{O}|}{p})$ times. Therefore, the overall *speedup* of Algorithm 2 over Algorithm 1 is at least,

$$\begin{aligned}
\text{speedup} =& \frac{O(d|\mathcal{O}|(\frac{\|\mathbf{X}\|_0}{q} + \frac{\|\mathbf{Z}\|_0}{p}))}{O(d(\|\mathbf{X}\|_0 + \|\mathbf{Z}\|_0 + |\mathcal{O}|))} \\
\geq& min\{O(\frac{|\mathcal{O}|}{q}), O(\frac{|\mathcal{O}|}{p})\} \\
=& O(|\mathcal{O}|/(q+p))
\end{aligned} \tag{19}$$

In application tasks, this can be at level of 10 to 100 such as collaborative filtering and link prediction. When $\|\mathbf{X}\|_0 + \|\mathbf{Z}\|_0$ is close to (or smaller than) $|\mathcal{O}|$ (datasets like Yahoo! Music, Tencent Weibo and Movielens-10M), our algorithm runs as fast as the algorithm of plain matrix factorization even though it uses extra features.

For the complexity of the parallel and efficient learning-to-match algorithm (PL2M) described in Algorithm 3, using K threads to run the algorithm, the computation cost for one round update is $O(\frac{1}{K}d(\|\mathbf{X}\|_0 + \|\mathbf{Z}\|_0 + |\mathcal{O}|))$. It is due to the fact that all parts of the algorithm are parallelized. This analysis does not consider the synchronization cost. In real world settings, we need to take synchronization cost into consideration, the corresponding time complexity becomes $O(\frac{1}{K}d((\|\mathbf{X}\|_0 + \|\mathbf{Z}\|_0)(1+\sigma) + |\mathcal{O}|))$, where σ denotes variance of computation costs by parallel tasks. Assume that we have K tasks and the time costs of the tasks are $T_1, T_2, \cdots T_K$. We define $\sigma = max_i(T_i)/Mean_i(T_i)$, since the training is delayed by the slowest task. To achieve maximum speedup, we need to schedule the tasks well such that the load of each task is average, which is always feasible when $|S|$, p, and q are large. Therefore, our algorithm can gain almost K times speedup.

In real world applications, there is a trade-off between the size of *parallel coordinate set* $|S|$ and the parameter σ, especially when different features have different levels of sparsity in the dataset. When we increase the size of *parallel coordinate set* $|S|$, we can divide the task into K threads in a more balanced way. On the other hand, η_{ks} will decrease as we increase $|S|$, making the update more conservative. Thus a parallel coordinate set needs to be chosen to balance convergence and acceleration. In fact, we need to empirically choose S such that each instance is covered by only a few nonzero features and the task size is large enough to run in a fairly balanced way.

In this paper, we fix $|S|$ and randomly partition elements from the feature indices to generate a set of disjoint subsets S in each round. We note that there can be more sophisticated scheduling strategies to select S, which is beyond the scope of this paper and can be an interesting topic for future research.

5. RELATED WORK

MF models [11] are arguably the most successful approach to learning-to-match. They have been applied to a wide range of real world problems, especially FMF models, which achieve state-of-the-art results, outperforming the models of MF in many different tasks, with different types of features used. In collaborative filtering, user feedback information (SVD++) [9], user attribute information [1], and product attribute information [23] are incorporated into models to further enhance the accuracies in prediction. In web search [25, 26], term vectors of queries and documents are utilized as features to significantly improve relevance ranking. FMF models also give the best results in link prediction in KDD Cup 2012 [14, 3, 4]. These works demonstrate the effectiveness of the learning-to-match models, but also create necessity for parallelization of the learning algorithms.

There has been much effort on parallelizing the process of plain matrix factorization. For example, Gemulla et al. [8] propose a method of distributed stochastic gradient descent for MF. Yu et al. [27] introduce a parallel coordinate descent algorithm for MF. An alternating least square method is proposed for MF as well [28]. Recently, Zhuang et al. [29] improve the efficiency of parallel stochastic gradient descent for MF by making a better scheduling of updates. Liu et al. [13] propose a distributed algorithm for nonnegative matrix factorization for web dyadic data analysis. The method of Probabilistic Latent Semantic Indexing is parallelized for Google news recommendation [7]. However, all the models on parallelizing plain matrix factorization replies on the fact that the rows and columns can be naturally separated and the parameters can be independently updated, and therefore cannot work on FMF due to the complex feature dependencies in the updating steps.

There is only a little work focusing on acceleration of coordinate descent for FMF. The most recent one scales up coordinate descent by making use of repeating patterns of features [19]. However, it is specialized for the least-squares loss and probit loss. Although the idea of avoiding repeated calculations is similar, our algorithm takes a completely different approach and can handle any general convex loss functions. Other algorithms of parallel coordinate descent, such as [6] and [15] cannot be directly applied to FMF, because it is difficult for them to handle complex feature

Table 1: Details of 4 Datasets

Dataset	$q \times p$	$\|\mathbf{X}\|_0$	$\|\mathbf{Z}\|_0$	$\|\mathcal{O}\|$	Task	Available Features
Yahoo! Music	$1M \times 0.6M$	$263M$	$1M$	$250M$	Collaborative Filtering	User Feedback [9], Taxonomy
Tencent Weibo	$2M \times 5K$	$55M$	$83K$	$93M$	Social Link Prediction	Social Network, User Profile, Taxonomy, Tag
Flickr	$0.65M \times 20K$	$451M$	$20K$	$39M$	Image Tagging	MAP, Sift Descriptors of Image
Movielens-10M	$69K \times 10K$	$10M$	$10K$	$9M$	Collaborative Filtering	User Feedback [9]

dependencies in FMF.

As general parallelization technique, the Hogwild! algorithm [16] can be applied to our problem. However, its time complexity is $O(d|\mathcal{O}|(\frac{\|\mathbf{X}\|_0}{q}+\frac{\|\mathbf{Z}\|_0}{p}))$, the same as Algorithm 1, due to the repeated calculations. Using the same number of threads, as analyzed in time complexity sections, it is theoretically $O(|\mathcal{O}|/(q+p))$ times slower than our parallel algorithm. In experiments, our parallel algorithm runs averagely about 5 times faster than Hogwild!. Another thread of related work is parallelization of coordinate descent algorithms. There have been studies on parallelizing coordinate descent for linear regression [2, 20, 21], other than matrix factorization [27]. The convergence of these algorithms depends on the spectrum of covariance matrix, which changes in each round in our learning setting (due to the changes in $\mathbf{U}$ and $\mathbf{V}$), and thus the algorithms cannot be directly applied to our problem. Our algorithm makes use of parallel update to minimize an upper bound re-estimated each round to ensure convergence, which can also be viewed as a kind of minorization-maximization algorithm [12].

6. EXPERIMENTS

In this section, we introduce our experimental results on several matching tasks using benchmark datasets. We first conduct comparison on accuracies between feature-based matrix factorization and plain matrix factorization. We then make comparisons on accuracies and efficiencies between our method of parallel learning-to-match and the baselines, including Hogwild! [16]. Finally, we conduct analysis on the efficiency of our parallel learning algorithm.

6.1 Datasets

Four datasets representing different types of learning-to-match tasks are chosen. Details of the datasets are summarized in Table 1.

The first dataset is Yahoo! Music Track1 [3] from the Yahoo! Music website. The dataset is among the largest public datasets for collaborative filtering. We use the official split of the dataset for experiments. As features, we use the implicit feedback of users [9] as well as the taxonomical information between the tracks, albums and artists, in addition to the indicators of users and tracks. Because it is an item rating dataset, we choose square loss as the loss function and use Root Mean Square Error (RMSE) as evaluation measure.

The second dataset is Tencent Weibo (microblog)[4], for social link prediction. The task is to predict a potential list of celebrities that a user will follow. The dataset is split into training and test data by time, with the test data further split into public and private sets for independent evaluations. We use the training set for learning and the public test set for evaluation. We use logistic loss as the loss function and MAP@K as evaluation metric, which is officially adopted in the KDD Cup competition [5]. The matrix data is extremely sparse, with only on average two positive links per user. Furthermore, about 70% of users in the test set have no following records in the training set. However, there are lots of additional information available, including social network and interaction (i.e., retweeting and commenting) records, profiles of users, categories of celebrities, and tags/keywords of users. The information is used as features of the task.

The third dataset is for automatic annotating images crawled from Flickr[6]. The dataset contains 0.65 million images and each image is associated with on average four tags. We select the $20,000$ most frequently occurring tags as the tag set. We randomly select $10,000$ images as test set and use the rest of images as training set. We use the bag-of-words vector of SIFT descriptors as features for images, and indicator vectors as features for tags. Logistic loss is chosen as the loss function. In testing, we generate a rank list of tags and use P@K(Precision at K) and MAP as evaluation metrics.

The fourth dataset is also for collaborative filtering, provided by Movielens[7]. We use the official split of dataset for experiments. This dataset is added because Hogwild! cannot run on Yahoo! Music dataset due to its high time complexity. In addition to the indicators of users and movies, the implicit feedbacks of users [9] are used as features. Similar to Yahoo! Music dataset, we choose square loss as the loss function and RMSE as evaluation metric.

6.2 Experiment Setting

We have implemented our parallel and efficient algorithm for learning-to-match (PL2M) using OpenMP[8]. The experiments are conducted on a machine with an Intel Xeon CPU E5-2680 (8 cores, supporting 16 threads at 2.70GHz, 128GB memory). We utilize up to 15 working threads and reserve one thread for scheduling.

We compare the performance of PL2M with those of serial algorithm for learning-to-match algorithm (denoted as L2M) and the Hogwild! algorithm [16]. To simplify the notations, here we use PL2M-T to refer to the parallel algorithm for learning-to-match with parallel set $|S| = T$ (e.g PL2M-5K means the parallel algorithm with $|S| = 5,000$). Hogwild! is the only one that can be directly applied to our problem as mentioned in Section 5. We have also implemented Hogwild! using OpenMP. *All matrix operations* mentioned in the algorithms take the advantages of data sparsity. PL2M, L2M , and Hogwild! share the same codes of elementary operations.

[3]http://kddcup.yahoo.com/datasets.php
[4]http://kddcup2012.org/c/kddcup2012-track1/data
[5]http://kddcup2012.org/c/kddcup2012-track1/details/Evaluation
[6]http://www.flickr.com
[7]http://www.movielens.org/
[8]http://www.openmp.org

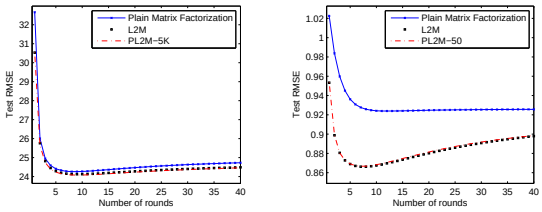

(a) on Yahoo! Music (b) on ML-10M

Figure 3: Test RMSE

Table 2: Results of Social Link Prediction (Tencent Weibo) in MAP

Setting	MAP@1	MAP@3	MAP@5
Popularity	22.54%	34.65%	38.28%
L2M-SNS	24.10%	36.56%	40.19%
PL2M-500-SNS	24.18%	36.65%	40.27%
L2M-ALL	25.44%	38.02%	41.63%
PL2M-500-ALL	25.52%	38.14%	41.75%

We empirically set $\lambda = 1, \alpha = 0.1$ and $d = 64$ for L2M and PL2M through all our experiments. To make fair comparison, the parameters of Hogwild! including learning rate, λ and α are tuned with cross validation on training set.

6.3 Usefulness of Features

We make comparison between FMF and MF to investigate the effectiveness of the features. We first compare FMF (by the algorithm of L2M) and MF in terms of test RMSE on the Yahoo! Music dataset in Figure 3(a). From the result, we can find that the FMF model converges faster and achieves better results than the MF model. This result is consistent with the result reported in [9, 18] and confirms the importance of using features in this problem. The test error first decreases but increases again in different rounds of training, indicating that training can stop at about 5 rounds.

The results of Tencent Weibo dataset are shown in Table 2. Since MF gives similar performance as the popularity based algorithm that only considers the popularity of each target node, and thus the result is not reported. Here the suffix ALL stands for the FMF model using all the available features shown in Table 1. We also evaluate the performance of the FMF model with only social network information, with suffix SNS. From Table 2, we can see that this dataset is extremely biased toward popular nodes. However, it is still possible to improve the results using social network information, and the auxiliary features can help to achieve the best performance. Note that PL2M-500-ALL has achieved the best result in Tencent Weibo dataset (actually our method is the same as the champion system on this dataset [4]).

The performance on Flickr test set is shown in Table 3. Because training and test images do not overlap, we cannot use MF to make prediction, and thus we adopt the use of popularity scores as a baseline. From the result, we can find that FMF can improve upon the popularity method, and assign relevant tags using image content features.

The test RMSE curves of different algorithms on Movielens-10M dataset are shown in Figure 3(b). From the result, we can find that the FMF model converges faster and achieves better results than the MF model. This result demonstrates the importance of using features in this problem. The test error first decreases but increases again in different rounds of training, which indicates that training can be stopped at about 7 rounds.

Table 3: Results of Image Tagging (Flickr) in Precision

Setting	MAP	P@1	P@3
Popularity	3.96%	4.63%	4.08%
L2M	7.18%	11.05%	8.76%
PL2M-50	7.59%	11.86%	9.19%

6.4 PL2M versus L2M

We make comparison between PL2M and L2M in terms of accuracy and efficiency.

As shown in Figure 3(a), Table 2∼3, and Figure 3(b), PL2M always gives comparable or even better test errors as L2M, indicating that PL2M makes no sacrifice on accuracy in the parallelization.

Figure 2 gives the training loss curves of PL2M and L2M. From the figure, we can observe that PL2M always converges following the lower bound given by L2M at the beginning. This is consistent with our theoretical result on convergence in Section 4. That is, if PL2M and L2M start with the same initial values, PL2M can perform at most as well as L2M.

The things get changed, however, as the training goes on. On Tencent Weibo dataset, PL2M-500 converges slightly better than L2M after 20 rounds. On Movielens-10M dataset, although the training loss curves of L2M, PL2M-500, and PL2M-50 are almost the same, the training loss of PL2M is lower in the end. This may due to the fact that the loss function is non-convex for $\mathbf{P}$ and $\mathbf{Q}$. After several updates, for example, 20 rounds, the values of $\mathbf{P}$ and $\mathbf{Q}$ are quite different so that the two methods finally converge to different local minimums.

From these figures, we can also observe that PL2M-50K converges slower than PL2M-5K in Yahoo! Music dataset, PL2M-5K converges slower than PL2M-500 on Tencent Weibo dataset, PL2M-500 converges slower than PL2M-50 on Flickr dateset. PL2M-500 converges a little slower than PL2M-50 on Movielens-10M dateset. These are consistent with our previous theoretical result that smaller $|S|$ leads to faster convergence.

6.5 PL2M versus Hogwild!

We make comparison between PL2M and Hogwild!. Since Hogwild! is based on stochastic gradient descent, some parameters such as learning rate need to be tuned. After fine tuning the parameters using cross validation on the training set for Hogwild!, including learning rate, l_2 coefficient λ, and l_1 coefficient α, we have obtained its performance in Table 4, including test error, running time of one round of training, and number of rounds needed to get the best test error. Both Hogwild! and PL2M use 8 threads.

We can see that the running time for each round of PL2M is much shorter than that of Hogwild!, while their test errors are similar. The difference of running time on Tencent Weibo is not as much as Movielens-10M and Flickr datasets because they have less features. This is consistent with our theoretical result about the time complexity in Section 3 and Section 4.

Furthermore, Hogwild! needs more training than PL2M to achieve its best test errors. For example, Hogwild! needs 20 rounds but PL2M needs only 7 rounds to achieve their best test errors in Movielens-10M. Therefore, the total running time for PL2M to get the best performance is much smaller

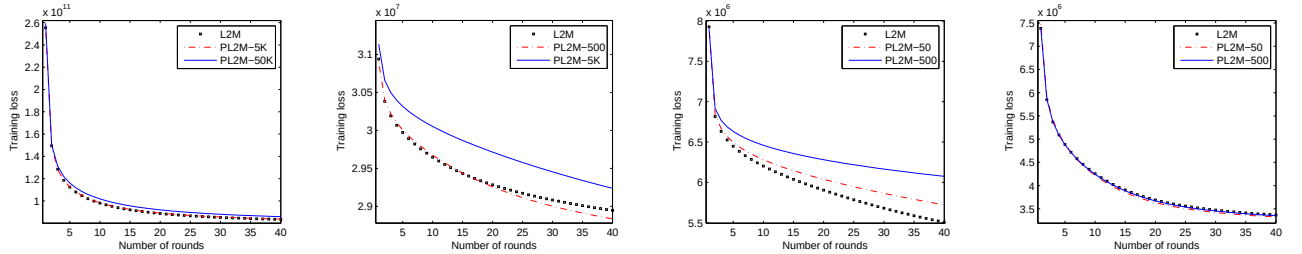

(a) TLC on Yahoo! Music (b) TLC on Tencent Weibo (c) TLC on Flickr (d) TLC on Movielens-10M

Figure 2: Training Loss Convergence(TLC) on Four Datasets

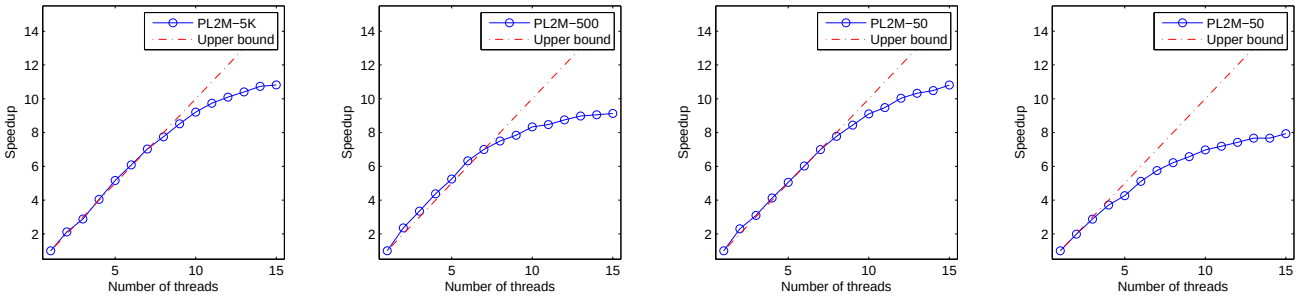

(a) on Yahoo! Music (b) on Tencent Weibo (c) on Flickr (d) on Movielens-10M

Figure 4: Efficiency Evaluation on Four Datasets (Speedup Curves)

Table 4: PL2M versus Hogwild! (8 threads)

Dataset	Method	sec/round	rounds	test error
Tencent Weibo	Hogwild!	104.3	14	24.96%
(MAP@1)	PL2M	70.1	5	25.52%
Flickr	Hogwild!	1117.0	59	7.59%
(MAP)	PL2M	155.0	33	7.59%
Movielens-10M	Hogwild!	162.2	20	0.8756
(RMSE)	PL2M	18.5	7	0.8666

than that of Hogwild!.

6.6 Scalability of PL2M

Finally, we evaluate the scalability of the parallel learning-to-match algorithm (PL2M). We test the average running time of PL2M-5K on Yahoo! Music dataset, PL2M-500 on Tencent Weibo dataset, PL2M-50 on Flickr dataset, PL2M-50 on Movielens-10M with varying numbers of threads and evaluate the improvement in efficiency.

As shown in Figure 4, the speedup curves are similar on Yahoo! Music, Tencent Weibo, and Flickr datasets, but the curve converges earlier on the Movielens-10M dataset. This is because that Movielens-10M is relatively smaller than the others and PL2M-50 runs really fast on Movielens-10M, which only needs about 18 seconds when 8 threads are used. Although the speedup gained by parallelization is not as much as that on other datasets, the parallel algorithm can also provide accelerations.

On the first 3 datasets, PL2M can achieve almost linear speedup with less than 8 threads, but the speedup gain slows down with more threads. We observe that the working threads are still fully occupied with more than 8. We conjecture that this turning point is due to the fact that the number of physical cores of the machine is only 8. From the results, we can find that PL2M is able to gain about 9 times speedup using 10 threads, confirming the scalability of the parallel algorithm.

In summary, the speedup gained by the parallel algorithm is significant, and thus it can easily handle hundreds millions of instances and features on a single machine.

7. CONCLUSION

We have proposed a parallel and efficient algorithm for learning-to-match, more specifically feature-based matrix factorization, a general and state-of-the-art approach. Our algorithm employs (1) iterative relaxations to solve the conflicts caused by parallel updates, with provable convergence guarantee on minimizing the original objective function, and (2) accelerate the computation by avoiding the repeated calculations caused by features, for any general convex loss functions. As a result, our algorithm can easily handle data with hundreds of millions of objects and features on a single machine. Extensive experimental results show that our algorithm is both effective and efficient when compared to the baselines.

As future work, we plan to (1) extend the algorithm to a distributed setting instead of the current multi-threading, (2) find better scheduling strategies for making parallel updates with a guaranteed bound of speedup, and (3) apply the technique developed in this paper to the parallelization of other learning methods, such as Markov Chain Monte Carlo (MCMC) learning methods for learning-to-match problem.

8. REFERENCES

[1] D. Agarwal and B.-C. Chen. Regression-based latent

factor models. KDD '09, pages 19–28, New York, NY, USA, 2009. ACM.

[2] J. Bradley, A. Kyrola, D. Bickson, and C. Guestrin. Parallel coordinate descent for l1-regularized loss minimization. In L. Getoor and T. Scheffer, editors, *Proceedings of the 28th International Conference on Machine Learning*, ICML '11, pages 321–328, New York, NY, USA, June 2011. ACM.

[3] K. Chen, T. Chen, G. Zheng, O. Jin, E. Yao, and Y. Yu. Collaborative personalized tweet recommendation. SIGIR '12, pages 661–670, New York, NY, USA, 2012. ACM.

[4] T. Chen, L. Tang, Q. Liu, D. Yang, S. Xie, X. Cao, C. Wu, E. Yao, Z. Liu, Z. Jiang, et al. Combining factorization model and additive forest for collaborative followee recommendation. *KDD CUP*, 2012.

[5] T. Chen, W. Zhang, Q. Lu, K. Chen, Z. Zheng, and Y. Yu. Svdfeature: A toolkit for feature-based collaborative filtering. *J. Mach. Learn. Res.*, 13(1):3619–3622, Dec. 2012.

[6] M. Collins, R. E. Schapire, and Y. Singer. Logistic regression, adaboost and bregman distances. *Machine Learning*, 48(1-3):253–285, 2002.

[7] A. S. Das, M. Datar, A. Garg, and S. Rajaram. Google news personalization: scalable online collaborative filtering. In *Proceedings of the 16th international conference on World Wide Web*, WWW '07, pages 271–280, New York, NY, USA, 2007. ACM.

[8] R. Gemulla, E. Nijkamp, P. J. Haas, and Y. Sismanis. Large-scale matrix factorization with distributed stochastic gradient descent. KDD '11, pages 69–77, New York, NY, USA, 2011. ACM.

[9] Y. Koren. Factorization meets the neighborhood: A multifaceted collaborative filtering model. KDD '08, pages 426–434, New York, NY, USA, 2008. ACM.

[10] Y. Koren. Collaborative filtering with temporal dynamics. volume 53, pages 89–97, New York, NY, USA, Apr. 2010. ACM.

[11] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, Aug. 2009.

[12] K. Lange, D. R. Hunter, and I. Yang. Optimization transfer using surrogate objective functions. *Journal of computational and graphical statistics*, 9(1):1–20, 2000.

[13] C. Liu, H.-c. Yang, J. Fan, L.-W. He, and Y.-M. Wang. Distributed nonnegative matrix factorization for web-scale dyadic data analysis on mapreduce. In *Proceedings of the 19th international conference on World wide web*, WWW '10, pages 681–690, New York, NY, USA, 2010. ACM.

[14] A. K. Menon and C. Elkan. Link prediction via matrix factorization. ECML PKDD'11, pages 437–452, Berlin, Heidelberg, 2011. Springer-Verlag.

[15] I. Mukherjee, K. Canini, R. Frongillo, and Y. Singer. Parallel boosting with momentum. In *Machine Learning and Knowledge Discovery in Databases*, pages 17–32. Springer, 2013.

[16] B. Recht, C. Re, S. Wright, and F. Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 693–701. 2011.

[17] S. Rendle. Factorization machines. In *Data Mining (ICDM), 2010 IEEE 10th International Conference on*, pages 995–1000. IEEE, 2010.

[18] S. Rendle. Factorization machines with libfm. *ACM Trans. Intell. Syst. Technol.*, 3(3):57:1–57:22, May 2012.

[19] S. Rendle. Scaling factorization machines to relational data. In *Proceedings of the VLDB Endowment*, volume 6, pages 337–348. VLDB Endowment, 2013.

[20] C. Scherrer, M. Halappanavar, A. Tewari, and D. Haglin. Scaling up coordinate descent algorithms for large ℓ_1 regularization problems. In ICML '12, pages 1407–1414, New York, NY, USA, July 2012. Omnipress.

[21] C. Scherrer, A. Tewari, M. Halappanavar, and D. Haglin. Feature clustering for accelerating parallel coordinate descent. In P. Bartlett, F. Pereira, C. Burges, L. Bottou, and K. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 28–36. 2012.

[22] J. Shang, T. Chen, H. Li, Z. Lu, and Y. Yu. A parallel and efficient algorithm for learning to match. In *Data Mining (ICDM), 2014 IEEE 14th International Conference on*. IEEE, 2014.

[23] D. H. Stern, R. Herbrich, and T. Graepel. Matchbox: large scale online bayesian recommendations. In *Proceedings of the 18th international conference on World wide web*, WWW '09, pages 111–120, New York, NY, USA, 2009. ACM.

[24] J. Weston, C. Wang, R. Weiss, and A. Berenzweig. Latent collaborative retrieval. In J. Langford and J. Pineau, editors, *Proceedings of the 29th International Conference on Machine Learning*, ICML '12, pages 9–16, New York, NY, USA, July 2012. Omnipress.

[25] W. Wu, H. Li, and J. Xu. Learning query and document similarities from click-through bipartite graph with metadata. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, WSDM '13, pages 687–696, New York, NY, USA, 2013. ACM.

[26] W. Wu, Z. Lu, and H. Li. Learning bilinear model for matching queries and documents. *J. Mach. Learn. Res.*, 14(1):2519–2548, Jan. 2013.

[27] H.-F. Yu, C.-J. Hsieh, S. Si, and I. Dhillon. Scalable coordinate descent approaches to parallel matrix factorization for recommender systems. ICDM '12, pages 765–774, Washington, DC, USA, 2012. IEEE Computer Society.

[28] Y. Zhou, D. Wilkinson, R. Schreiber, and R. Pan. Large-scale parallel collaborative filtering for the netflix prize. In *Proceedings of the 4th international conference on Algorithmic Aspects in Information and Management*, AAIM '08, pages 337–348, Berlin, Heidelberg, 2008. Springer-Verlag.

[29] Y. Zhuang, W.-S. Chin, Y.-C. Juan, and C.-J. Lin. A fast parallel sgd for matrix factorization in shared memory systems. RecSys '13, pages 249–256, New York, NY, USA, 2013. ACM.

[30] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.